Optimization for Speculative Execution of Multiple Jobs in a MapReduce-like Cluster

Huanle XU, Wing Cheong LAU
Department of Information Engineering, The Chinese University of Hong Kong
{xh112, wclau}@ie.cuhk.edu.hk

arXiv:1406.0609v3 [cs.DC] 5 Jan 2015

***Abstract*—Nowadays, a computing cluster in a typical data center can easily consist of hundreds of thousands of commodity servers, making component/ machine failures the norm rather than exception. A parallel processing job can be delayed substantially as long as one of its many tasks is being assigned to a failing machine. To tackle this so-called straggler problem, most parallel processing frameworks such as MapReduce have adopted various strategies under which the system may speculatively launch additional copies of the same task if its progress is abnormally slow or simply because extra idling resource is available. In this paper, we focus on the design of speculative execution schemes for a parallel processing cluster under different loading conditions. For the lightly loaded case, we analyze and propose two optimization-based schemes, namely, the Smart Cloning Algorithm (SCA) which is based on maximizing the job utility and the Straggler Detection Algorithm (SDA) which minimizes the overall resource consumption of a job. We also derive the workload threshold under which SCA or SDA should be used for speculative execution. Our simulation results show both SCA and SDA can reduce the job *flowtime* by nearly 60% comparing to the speculative execution strategy of Microsoft Mantri. For the heavily loaded case, we propose the Enhanced Speculative Execution (ESE) algorithm which is an extension of the Microsoft Mantri scheme. We show that the ESE algorithm can beat the Mantri baseline scheme by 18% in terms of job *flowtime* while consuming the same amount of resource.**

***Index Terms*—Job scheduling, speculative execution, cloning, straggler detection, optimization**

I. Introduction

Empirical performance studies of large-scale computing clusters have indicated that the completion time of a job [3] is often significantly and unnecessarily prolonged by one or a few so-called "straggler" (or outlier) tasks, i.e. tasks which are unfortunately assigned to either a failing or overloaded computing node within a cluster. As such, recent parallel processing frameworks such as the MapReduce system or its many variants have adopted various preventive or reactive straggler-handling strategies under which the system may automatically launch extra (backup) copies of a task on alternative machines in a judicious manner. Unfortunately, most of the existing speculative execution schemes are based on simple heuristics. In particular, there are two main classes of speculative execution strategies, namely, the Cloning [2] approach and Straggler-Detection-based one [3], [7], [9], [10], [13], [18]. Under the Cloning approach, extra copies of a task are scheduled in parallel with the initial task as long as the resource consumption of the task is expected to be low and there is system resource available. For the Straggler-Detection-based approach, the progress of each task is monitored by the system and backup copies are launched when a straggler is detected. As one may expect, the cloning-based strategy is only suitable for a lightly loaded cluster as it launches the clones in a greedy, indiscriminately fashion. On the other hand, the straggler-detection based strategy is applicable to both lightly-loaded and heavily-loaded regimes but at the expense of extra system instrumentation and performance overhead. The situation is particularly challenging when the progress of a large number of tasks have to be tracked.

In this paper, we take a more systematic, optimization-based approach for the design and analysis of speculative execution schemes. In particular, we have made the following technical contributions:

- After reviewing related work in Section II, we introduce the system model in Section III and derive the cut-off workload threshold between the lightly-loaded and heavily-loaded operating regimes of a computing cluster. Based on this workload threshold, the applicability of a speculative execution strategy can be analyzed for different operating regimes.
- In Section IV, we introduce a generalized cloning-based framework which can jointly optimize the job utility with resource consumption when the cluster is lightly loaded. We also present a specific Smart Cloning Algorithm (SCA) based on this framework.
- In Section V, we consider the optimal online-scheduling framework and design the Straggler Detection Algorithm (SDA) which launches an optimal number of extra copies for a straggler task in an on-demand basis, i.e. only after the straggler has been detected. In particular, we show that SDA can minimize overall resource consumption of the arriving jobs in a lightly loaded cluster.
- In Section VI, we propose the Enhanced Speculative Execution (ESE) algorithm for a heavily loaded cluster by extending the speculative execution strategy of Microsoft Mantri [3]. We demonstrate that ESE can improve the job completion time of Mantri while consuming the same amount of resource. We also summarize our findings and conclude the paper in Section VII.

II. Related work

Several speculative execution strategies have been proposed in the literature for the MapReduce system and its variants or derivatives. The initial Google MapReduce system only begins

to launch backup tasks when a job is close to completion. It has been shown that speculative execution can decrease the job execution time by 44% [9]. The speculative execution strategies in the initial versions of Hadoop [1] and Microsoft Dryad [10] closely follow that of the Google MapReduce system. However, the research group from Berkeley presents a new strategy called LATE (*Longest Approximate Time to End*) [18] in the Hadoop-0.21 implementation. It monitors the progress rate of each task and estimates their remaining time. Tasks with progress rate below certain threshold (*slowTaskTherehold*) are chosen as backup candidates and the one with the longest remaining time is given the highest priority. The system also imposes a limit on the maximum number of backup tasks in the cluster*speculativeCap*. Microsoft Mantri [3] proposes a new speculative execution strategy for Dryad in which the system estimates the remaining time to finish, t_{rem}, for each task and predicts the required execution time of a relaunched copy of the task, t_{new}. Once a computing node becomes available, the Mantri system makes a decision on whether to launch a backup task based on the statistics of t_{rem} and t_{new}. Specifically, a duplicate is scheduled if $\mathbb{P}(t_{rem} > 2 * t_{new}) > \delta$ is satisfied where the default value of $\delta = .25$. Hence, Mantri schedules a duplicate only if the total resource consumption is expected to decrease. Mantri also may terminate a task which shows an excessively large remaining-time-to-finish.

To accurately and promptly identify stragglers, [7] proposes a Smart Speculative Execution strategy and [13] presents an Enhanced Self-Adaptive MapReduce Scheduling Algorithm respectively. The main ideas of [7] include: i) use exponentially weighted moving average to predict process speed and compute the remaining time of a task and ii) determine which task to backup based on the load of a cluster using a cost-benefit model. Recently, [2] proposes to mitigate the straggler problem by cloning every small job and avoid the extra delay caused by the straggler monitoring/ detection process. When the majority of the jobs in the system are small, the cloned copies only consume a small amount of additional resource.

III. System Model

Assume a set of jobs $J = \{J_1, J_2, \cdots\}$ arriving at a computing cluster at a rate of λ jobs per unit time. Different jobs may run different applications and a particular job J_i which arrives at the cluster at time a_i consists of m_i tasks. This cluster has M computing nodes (machine) and each computing node can only hold one task at any time. For simplicity, we assume this cluster is homogeneous in the sense that all the nodes are identical. Further, we assume that the execution time (i.e. the time between the task is launched and the task is finished) of each task of J_i without any speculative execution follows the same distribution, i.e., $x_j^i \sim F_i(t) = Pr(x_j^i < t)$ for $1 \leq j \leq m_i$. We assume the execution time distribution information can be estimated for each job according to prior trace data of the application it runs and the size of the input data to be processed. Upon arrival, each job joins a queue in the master-node of the cluster, waiting to be scheduled for execution according to some priorities to be determined in the following sections.

Here, we define the job *flowtime* which is an important metric we capture as below.

Definition 1. *The flowtime of a job J is $flow(J_i) = f(J_i) - a(J_i)$, where $f(J_i)$ and $a(J_i)$ denote the finish (completion) time and arrive time respectively.*

If a task j runs t_j units of time on a computing node, then it consumes $\gamma * t_j$ units of resource on this node where γ is a constant number. We ignore the resource consumption of an idle machine. For ease of description, we often interchange the two notations in this paper, namely, machine and computing node.

A. Speculative execution under different operating regimes

The cloning-based strategy for speculative execution schedules extra copies of a task in parallel with the initial task as long as the resource consumption of the task is expected to be low and there is system resource available. Only the result of one which finishes first among all the copies is used for the subsequent computation. Cloning does not incur any monitoring overhead. Nevertheless, cloning consumes a large amount of resource and can easily block the scheduling of subsequent jobs when the cluster workload is heavy.

On the other hand, the Straggler-Detection-based approach makes a speculative copy for the task after a straggler is detected. In this approach, the scheduler needs to monitor the progress for each task. However, the monitoring incurs extra system instrumentation and performance overhead as discussed in [4]. The situation is particularly challenging when the progress of a large number of tasks have to be tracked. To make things even worse, it's always difficult to detect a straggler for small jobs as they usually complete their work in a very short period [2].

As one may expect, the cloning-based strategy is only suitable for a lightly loaded cluster as it launches the clones in a greedy, indiscriminately fashion. On the other hand, the straggler-detection based strategy is applicable to both lightly-loaded and heavily-loaded regimes. Hence, there exists a cutoff threshold to separate the cluster workload into these two operating regimes. When the workload is below this threshold, the cloning-based strategy can obtain a good performance in terms of job *flowtime*. Conversely, when the workload exceeds this threshold, only the Straggler-Detection-based strategy can help to improve the cluster performance.

B. Deriving the cutoff threshold for different regimes

In this subsection, we derive the cutoff workload threshold λ^U which allows us to separate our subsequent analysis into the lightly loaded vs. heavily loaded regimes. We assume that the random variables m_i and x_j^i are independent from each other for all i and j. To simplify the analysis, we focus on the task delay in the cluster instead of job delay. Denote by λ_t the task arrival rate to the whole cluster. Thus, $\lambda_t = \mathbb{E}[m_i] \cdot \lambda$. We approximate the task arrival as a Poisson process with rate

λ_t. Further, denote by λ_t^m the average task arrival rate to each single machine which is given by $\lambda_t^m = \frac{\lambda_t}{M}$.

We model the task service process of each computing node (machine) as a $M/G/1$ queue. Applying the result of [8], we get the average delay of each task without speculative execution made in the following equation

$$W_t = \frac{\lambda_t^m \mathbb{E}[s^2]}{2(1-\lambda_t^m \mathbb{E}[s])} + \mathbb{E}[s] \tag{1}$$

where $\mathbb{E}[s]$ is the average duration of all the tasks in the system without speculative execution implemented.

We proceed to derive the expression of the task delay for the cloning-based strategy. Different from [2] where the cloning is done for small jobs only, here we consider a more general cloning scheme in which the small jobs are not distinguished from the big ones. In this scheme, each task should at least make two copies. Otherwise, the task which does not have any extra copy may delay the completion of the entire job.

We illustrate an example in which $F_i(t)$ (for all i) follows the Pareto Distribution as follows:

$$F_i(t) = \begin{cases} 1-(\frac{\mu}{t})^\alpha & for\ t \geq \mu \\ 0 & otherwise \end{cases}$$

Assume r (≥ 2) copies are launched for a particular task v. Then the expected duration for task v is $\mathbb{E}[s^{'}] = \frac{\mu r \alpha}{r\alpha - 1}$. Thus, $\frac{\mathbb{E}[s^{'}]}{\mathbb{E}[s]} = \frac{\alpha-1}{\alpha-1/r} > \frac{\alpha-1}{\alpha}$. This also gives a lower-bound of the performance improvement of cloning regardless of the number of extra copies to be made for each task.

Denote by $\mathbb{E}[r]$ the average number of copies each task makes. Hence, $\mathbb{E}[r] \geq 2$. Further define $\mathbb{E}[s^c]$ and λ_t^c as the average task duration and equivalent task arrival rate to each machine respectively after the cloning is made.

The first constraint for cloning is that it must not overload the system, i.e. the long-term system utilization of the cluster should be less than 1. Thus, the following inequality holds:

$$\lambda_t^c \cdot \mathbb{E}[s^c] < 1 \tag{2}$$

By considering the constraint in Eq.(2), and the fact that $\lambda_t^c = \mathbb{E}[r] \cdot \lambda_t^m$, we have:

Theorem 1. *The condition $\lambda \cdot \mathbb{E}[m_i]\mathbb{E}[s] \cdot \frac{4(\alpha-1)}{2\alpha-1} < M$ is necessary to guarantee that the cloning does not overload the system.*

Proof. Refer to the technical report [17]. □

However, the efficiency of cloning is not guaranteed by Theorem 1. An efficient cloning strategy should have a smaller task delay than a strategy which does not make speculative execution. This argument must also hold when each task has only two copies. Denote by W_t^c the average task delay when each task has two copies. For convenience, we define $\omega = \frac{\lambda \cdot \mathbb{E}[m_i]\mathbb{E}[s]}{M}$. Thus,

$$W_t^c = \mathbb{E}[s] \cdot \frac{\omega \cdot \frac{(\alpha-1)(1-4\alpha^2+4\alpha)}{\alpha(2\alpha-1)} + 2(\alpha-1)}{2\alpha - 1 - 4\omega(\alpha-1)} \tag{3}$$

and

$$W_t^c < W_t \tag{4}$$

Combine (1), (3) and (4), we can derive the upper bound ω^U for ω. Hence, the cutoff threshold is determined by the following equation:

$$\lambda^U = \frac{\omega^U M}{\mathbb{E}[m_i]\mathbb{E}[s]} \tag{5}$$

In the following sections, we continue to introduce the cloning-based strategy and Straggler-Detection-based approaches under two different workload regimes.

IV. Optimal Cloning in the Lightly Loaded Regime

In the lightly loaded cluster, i.e., $\lambda < \lambda^U$, we first apply the generalized cloning-based scheme to improve the job performance.

We consider that time is slotted and the scheduling decisions are made at the beginning of each time slot. Assume $J_i \in J$ consisting of m_i tasks which are from the set $\Phi_i = \{\delta_1^i, \delta_2^i, \cdots, \delta_{m_i}^i\}$ and δ_j^i is scheduled at time slot w_j^i. Denote by w_i the scheduling time of job i. Hence, w_j^i and w_i satisfy the following constraints:

$$w_j^i \in \{0, 1, 2\cdots,\} \quad and \quad w_j^i \geq a_i \quad \forall i;\ j \tag{6}$$

$$w_i = \min\{w_1^i, w_2^i, \cdots, w_{m_i}^i\} \quad \forall i \tag{7}$$

Each task of Φ_i can maintain different number of duplicates as the tasks in the same job may be scheduled at different time slots depending on server availability. Denote by c_j^i the number of copies made for task δ_j^i. Further let $t_{j,k}^i$ define the duration of the kth clone for task δ_j^i. We assume $t_{j,k}^i$ follows the same distribution as x_j^i and all the $t_{j,k}^i$ are i.i.d random variables for $1 \leq k \leq c_j^i$. Define t_j^i as the duration of task δ_j^i and t_i as the *flowtime* of job J_i respectively. Then the following two equations hold:

$$t_j^i = \min\{t_{j,1}^i, t_{j,2}^i, \cdots, t_{j,c_i}^i\} \qquad \forall i; 1 \leq j \leq m_i \tag{8}$$

$$t_i = \max\{t_1^i + w_1^i, t_2^i + w_2^i, \cdots, t_{m_i}^i + w_{m_i}^i\} - a_i \qquad \forall i \tag{9}$$

Equation (8) states that as soon as one copy of task δ_j^i finishes, the task completes. Equation (9) describes the job *flowtime*.

We define a utility for each job which is a function of job *flowtime* and the number of tasks it maintains. The formulation (P1) is as follows:

$$\begin{aligned} \max_{c_j^i, w_j^i} \quad & \sum_{i=1} U\left(\mathbb{E}[t_i], m_i\right) - \gamma \cdot \sum_{i=1} \sum_{j=1}^{m_i} c_j^i \cdot \mathbb{E}[t_j^i] \\ s.t. \quad & \sum_{w_i \leq l} \sum_{w_j^i + t_j^i > l} c_j^i \leq M \quad \forall l \\ & 1 \leq c_j^i \leq r \quad \forall i; 1 \leq j \leq m_i \\ & (6), (8), (9) \end{aligned}$$

In this formulation, the utility of job J_i, $U\left(\mathbb{E}[t^i], m_i\right)$ is a strictly concave and differentiable function of $\mathbb{E}[t^i]$ and m_i. Our objective is to maximize the total utility of all the

jobs in the cluster while keeping a low resource consumption level. The first constraint states that the total number of tasks including all task copies at any time slot is no more than M and the second constraint states that each individual task can at most maintain r copies in the cluster.

A. Solving P1 through approximation

P1 is an online stochastic optimization problem and the scheduling decisions should be made without knowing the information of future jobs. In Equation (9), the tasks of the same job can be scheduled in different time slots. Hence, it is not easy to express the job *flowtime* in terms of the distribution function and thus makes P1 difficult to solve.

Due to the fact that the cluster is lightly loaded, there is a large room for making clones for all the jobs most of the time. Thus, we solve another optimization problem (P2) as a relaxation for P1 when system resource is available. In P2, all the tasks of the same job are scheduled together and maintain the same number of copies. In this way, we can simplify the modeling of the job *flowtime*.

Define $\chi(l)$ as the job set which contains all unscheduled jobs at time slot l. Assume $\chi(l) = \{J_{l_1}, J_{l_2}, \cdots\}$. If there is enough idling servers to schedule the jobs in the cluster at the beginning of time slot l, i.e., $\sum_i m_{l_i} < N(l)$ where $N(l)$ is number of available machines, we solve P2 to determine the number of copies for each task in J_{l_i} as below:

$$\begin{aligned}
\max_{c_{l_i}} \quad & \sum_{i=1} U\left(\mathbb{E}[t_{l_i}], m_{l_i}\right) - \gamma \cdot \sum_{i=1} \sum_{j=1}^{m_{l_i}} c_{l_i} \cdot \mathbb{E}[t_j^{l_i}] \\
s.t. \quad & \sum_i m_{l_i} \cdot c_{l_i} \le N(l) \\
& 1 \le c_{l_i} \le r \quad \forall i \\
& t_j^{l_i} = \min\{t_{j,1}^{l_i}, t_{j,2}^{l_i}, \cdots, t_{j,c_{l_i}}^{l_i}\} \qquad \forall i; 1 \le j \le m_{l_i} \\
& t_{l_i} = \max\{t_1^{l_i}, t_2^{l_i}, \cdots, t_{m_i}^{l_i}\} + l - a_{l_i} \qquad \forall i
\end{aligned}$$

Here, c_{l_i} is the number of duplicates assigned to each task in job J_{l_i}. Solving P2 is much easier and it only depends on the current information. Define $H_j^{l_i}(t)$ as the cumulative distribution function of $t_j^{l_i}$ and we have:

$$H_j^{l_i}(t) = 1 - (1 - F_{l_i}(t))^{c_{l_i}} \tag{10}$$

Let $d_{l_i} \triangleq \max\{t_1^{l_i}, t_2^{l_i}, \cdots, t_{m_i}^{l_i}\}$ and the distribution function of d_{l_i} is given by:

$$W_{l_i}(t) = Pr(d_{l_i} < t) = \prod_{j=1}^{m_{l_i}} H_j^{l_i}(t) = (1-(1-F_{l_i}(t))^{c_{l_i}})^{m_{l_i}} \tag{11}$$

Further, $\mathbb{E}[t_{l_i}] = \mathbb{E}[d_{l_i}] + l - a_i$ and we get

$$\mathbb{E}[d_{l_i}] = \int_0^\infty t \cdot d(W_{l_i}(t)) = \int_0^\infty (1 - W_{l_i}(t))dt. \tag{12}$$

Similarly, $\mathbb{E}[t_j^{l_i}] = \int_0^\infty (1 - H_j^{l_i}(t))dt$, which yields:

$$\sum_{j=1}^{m_{l_i}} c_{l_i} \cdot \mathbb{E}[t_j^{l_i}] = m_{l_i} c_{l_i} \int_0^\infty (1 - F_i(t))^{c_{l_i}} dt \tag{13}$$

Lemma 1. *$c_{l_i} \int_0^\infty (1 - F_i(t))^{c_{l_i}} dt$ is a convex function of c_{l_i} when provided $\ln(1 - F_i(t))$ is a convex function of t.*

Proof. Refer to [17]. □

In the same way, $\mathbb{E}[t_{l_i}]$ is also a convex function of c_{l_i} which decreases as c_{l_i} increases.

In traditional scheduling problems, minimizing the job *flowtime* is a common objective. Thus, we consider to minimize the summation of job *flowtime* and resource consumption as a special example, i.e.,

$$U\left(\mathbb{E}[t_{l_i}], m_{l_i}\right) = -\mathbb{E}[t_{l_i}].$$

Observe that the first two constraints in P2 are linear. Hence, we can adopt the convex optimization technique to solve P2. The Lagrangian Dual problem of P2 is given by:

$$\min_{\nu, \xi_{l_i}, h_{l_i}} D(\nu, \xi_{l_i}, h_{l_i}) \tag{14}$$

$$D(\nu, \xi_{l_i}, h_{l_i}) = \max_{c_{l_i}} f(\nu, \xi_{l_i}, h_{l_i}, c_{l_i}) \tag{15}$$

$$\begin{aligned}
f(\nu, \xi_{l_i}, h_{l_i}, c_{l_i}) = & -\sum_{i=1} \mathbb{E}[t_{l_i}] - \gamma \cdot \sum_{i=1} m_{l_i} c_{l_i} \cdot \mathbb{E}[t_j^{l_i}] \\
& - \nu(\sum_{i=1} m_{l_i} c_{l_i} - N(l)) - \sum_i \xi_{l_i}(c_{l_i} - r) \\
& - \sum_{i=1} h_{l_i} \cdot (1 - c_{l_i})
\end{aligned} \tag{16}$$

where ν, ξ_{l_i}, h_{l_i} are nonnegative multipliers. Applying the result of convex optimization, we conclude that there is no duality gap between P2 and D.

We adopt the gradient projection algorithm to get the optimal solution of D. Define the vector $c(l) = (c_{l_1}, c_{l_2}, \cdots)$ and the algorithm is outlined below:

- Initialize $c_{l_i}^0 = 1$ for all i, $\nu^0 = 0.1$, $\xi_{l_i}^0 = 0.1$, $h_{l_i}^0 = 0.1$.
- $c_{l_i}^{k+1} = \arg\max_{c_{l_i}} f(\nu^k, \xi_{l_i}^k, h_{l_i}^k, c_{l_i})$;
- $v^{k+1} = v^k + \eta_1[\sum_{i=1} m_{l_i} c_{l_i}^{k+1} - N(l)]_{v^k}^+$;
- $\xi_{l_i}^{k+1} = \xi_{l_i}^k + \eta_2[c_{l_i}^{k+1} - r]_{\xi_{l_i}^k}^+$;
- $h_{l_i}^{k+1} = h_{l_i}^k + \eta_3[1 - c_{l_i}^{k+1}]_{h_{l_i}^k}^+$;
- if $|c^{k+1}(l) - c^{k+1}(l)| < \epsilon$, the gradient algorithm terminates.

where $\{c, v, \xi, h\}_*^k$ denote the values of the corresponding parameters during the kth iteration of the algorithm.

Next, we proceed to prove the convergence of this algorithm by adopting the method of Lyapunov stability theory [11]. We first define the following Lyapunov function:

$$\begin{aligned}
V(\nu, \xi_{l_i}, h_{l_i}) = & \int_{\nu^*}^{\nu} \frac{(\xi - \nu^*)}{\eta_1} d\xi + \sum_i \int_{\xi_{l_i}^*}^{\xi_{l_i}} \frac{(\xi - \xi_{l_i}^*)}{\eta_2} d\xi \\
& + \sum_i \int_{h_{l_i}^*}^{h_{l_i}} \frac{(\xi - h_{l_i}^*)}{\eta_3} d\xi
\end{aligned} \tag{17}$$

where ν^*, $\xi_{l_i}^*$ and $h_{l_i}^*$ is the optimal solution of D. It can be shown that $V(\nu, \xi_{l_i}, h_{l_i})$ is positive definite. (Refer to [17] for the details of the proof.)

Denote by $\dot{V}$ the derivative of $V(\nu, \xi_{l_i}, h_{l_i})$ with respect to time. With the following lemma, we get $\dot{V} \leq 0$.

Lemma 2. *The subgradient of $D(\nu, \xi_{l_i}, h_{l_i})$ at ν is given by:*

$$\frac{\partial D}{\partial \nu} = N(l) - \sum_{i=1} m_{l_i} \tilde{c}_{l_i}$$

Similarly, the subgradient of $D(\nu, \xi_{l_i}, h_{l_i})$ and ξ_{l_i}, h_{l_i} are:

$$\frac{\partial D}{\partial \xi_{l_i}} = r - \tilde{c}_{l_i}; \quad \frac{\partial D}{\partial h_{l_i}} = \tilde{c}_{l_i} - 1$$

respectively where $\tilde{c}_{l_i}$ minimizes $f(\nu, \xi_{l_i}, h_{l_i}, c_{l_i})$.

Proof. Refer to [17] for details. □

Theorem 2. *When the step size η_1, η_2, η_3 are positive, the above gradient projection algorithm can converge to the global optimal.*

Proof. First, the trajectories of V converge to the set $S = \{(\nu, \xi_{l_i}, h_{l_i}) | \dot{V}(\nu, \xi_{l_i}, h_{l_i}) = 0\}$ according to the Lasalle's Principle [11]. Then, based on the proof of $\dot{V} \leq 0$ in [17], for all $(\nu, \xi_{l_i}, h_{l_i}) \in S$, $D(\nu, \xi_{l_i}, h_{l_i}) = D^*$ which is the optimal value. Thus, all the elements in S must be global optimal solutions. Hence, we conclude that the gradient projection algorithm converges to the global optimal. □

Not only that the algorithm is guaranteed to converge to the global optimal, we can further tune its parameters to speed up the convergence in an actual implementation. Fig. 1 depicts the results of a matlab-based simulation experiment to demonstrate the fast convergence rate of the gradient projection algorithm. In this experiment, we assume that, in a particular time slot l, there are 4 jobs waiting to be scheduled, i.e., $\chi(l) = \{J_{l_1}, J_{l_2}, J_{l_3}, J_{l_4}\}$. The cluster has 100 available machines and the number of tasks for each job are 10, 20, 5 and 10 respectively. Assume $F_i(t)$ to follow the Pareto Distribution where $F_i(t) = 1 - (\frac{\mu_i}{t})^2$ for $t \geq \mu_i$ and $\mu_1 = 1, \mu_2 = 2, \mu_3 = 1, \mu_4 = 2$ and the number of copies for each task is given by $r = 8$. We tune the parameters η_1, η_2, η_3 to be 0.2, 0.3, 0.4 respectively. As shown in Fig. 1, this algorithm can converge very fast to the optimal solution.

B. The design of the Smart Cloning Algorithm (SCA)

Following the analysis in subsection IV-A, there exists a case where there is no space for cloning the jobs at the beginning of a particular time slot. In this scenario, it does not make sense to solve P2. Instead, we adopt a smallest remaining workload first scheme. It is well known in scheduling literature that the Shortest Remaining Processing Time (SRPT) scheduler is optimal for overall flowtime on a single machine where there is one task per job. The SRPT-based approach has been adopted widely for the scheduling in a parallel system as presented in [5] [6], [12] [14] [15] [16], [19]. Based on P2 and SRPT, we propose the Smart Cloning Algorithm (SCA) below.

SCA consists of two separate parts. At the beginning of each time slot, we first schedule the remaining tasks of the unfinished jobs and then check whether the condition $\sum_i m_{l_i} < N(l)$ is satisfied. If the condition is satisfied, we solve P2 to determine the number of clones for each task. Otherwise, i.e. $\sum_i m_{l_i} \geq N(l)$, we sort $\chi(l)$ according to the increasing order of the workload in each J_{l_i} and schedule the jobs based on this order and only one copy of each task is created. Notice that the resultant workload is the product of m_{l_i} and $\mathbb{E}[x_j^i]$. The corresponding pseudo-code is given in Algorithm 1 as below.

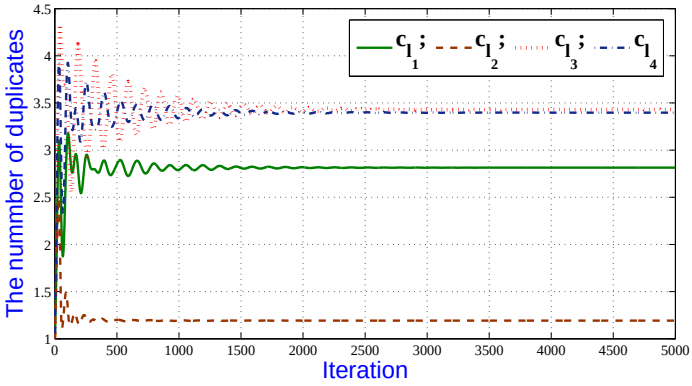

Fig. 1. The convergency performance of the gradient projection algorithm. The number of duplicates for the tasks in each job iterates and can converge to the optimal value. c_{l_i} represents the number of duplicates for each task in job l_i where $i = 1, 2, 3, 4$.

Input: The jobs in the cluster associated with their running status at time slot l;
Output: Scheduling decisions for time slot l.

```
schedule the unassigned tasks of the running jobs in the cluster with the fewest remaining first;
update N(l) and χ(l);
if N(l) == 0 then
    return;
end
if Σ_i m_{l_i} < N(l) then
    solve P2 and assign duplicates of the tasks in J_{l_i} based on the optimization result;
else
    for Job J_{l_i} in χ(l) do
        assign only one copy for each task of J_{l_i};
        update N(l);
        if N(l) == 0 then
            return;
        end
    end
end
return;
```

Algorithm 1: Smart Cloning Algorithm

C. Performance evaluation for SCA

We run a Matlab-based simulation to evaluate the performance of SCA. In the simulation, the jobs arrive at the computing cluster following a Poisson process with rate $\lambda = 6$.

There are $M = 3000$ machines in the cluster. The number of tasks in each job is uniformly distributed from 1 to 100. Within a job, the duration of each task follows a common Pareto distribution with a heavy tail order of 2. The expected task duration for different jobs follows a uniform distribution of 1 to 4 units of time. The resource consumption parameter γ is set to be 0.01. The simulation is run for 1500 units of time and repeated using 3 different seeds.

We use the speculative execution strategy of Microsoft Mantri as the baseline and compare it with the proposed SCA. We use job *flowtime* and job resource consumption as the performance metrics for comparison. Fig. 2 depicts the cumulative density function of job *flowtime* and resource consumption for nearly 27000 jobs where the solid line represents the results of SCA. Observe from the figure that the average job *flowtime* reduces by 60% in our algorithm compared to the baseline. It is worthnoting that, under SCA, more than 80% (90%) of jobs can finish within 6 (9) time units respectively. As a comparison, about 80% (90%) of jobs can finish within 17 (25) time units for the Mantri algorithm. However, our algorithm consumes more resource than the Mantri one. It indicates that 80% of jobs consume less than 1.5 units resource in the Mantri baseline while 80% of jobs consume less than 2 units resource in SCA. Recall that the objective of Smart Cloning Algorithm is to maximize the difference of job utility and resource consumption. To further make a fair comparison between these two schemes, we consider the job utility minus total resource consumption as an additional performance metric. Our simulation results show that SCA can beat the baseline considerably in terms of this metric.

V. Design of Optimal Straggler-Detection-Based Scheme for the Lightly Loaded Regime

As analyzed in Section III, both the cloning and straggler detection approach can improve the system performance in the lightly loaded regime. However, it is still unknown which one has a better performance. In this section, we formulate a different model to design speculative executions based on individual-task-progress monitoring. Different from the previous Straggler-Detection-based approaches which only duplicate one copy at most for each straggler, this model can automatically determines the optimal number of duplicates when a straggler is detected.

We use the same notations as in Section IV except for c_i. Based on the monitoring result, the scheduler only begins to make speculative copies for task δ_j^i if a straggler is detected, i.e. when the progress of δ_j^i substantially falls behind others. To be more specific, a task is declared to be a straggler if its estimated remaining time to finish (t_{rem}) is greater than $\sigma_i \cdot \mathbb{E}[t_{new}]$ where t_{new} is the expected execution time of a new copy. In this model, different tasks in the same job can maintain different number of duplicates based on their task progress. We define $c_j^i(t)$ as the number of copies running for task δ_j^i at time $t \geq 0$. Under this setup, σ_i is a parameter to be optimized for all i. Intuitively, if σ_i is too small, a lot of running tasks will be characterized as stragglers. This can incur a large number of duplicates and consume a lot of resource in the cluster. On the other hand, if σ_i is set to be too large, many jobs will be delayed due to a small number of slow-progressings tasks as the speculative copies of tasks are not launched.

To model the monitoring progress, we assume that the scheduler can detect the straggler for task δ_j^i after it completes a portion of the work which is s_i. Here, s_i is a constant number related to job J_i. Thus, $c_j^i(t)$ needs to satisfy the following two constraints:

$$c_j^i(t) = 1 \quad if \ w_j^i \leq t \leq s_i \cdot t_{j,1}^i + w_j^i \tag{18}$$

$$c_j^i(t) = 1 \quad if \ (1-s_i) \cdot t_{j,1}^i < \sigma_i \cdot \mathbb{E}[x_j^i] \ \& \ w_j^i \leq t \leq t_{j,1}^i + w_j^i \tag{19}$$

Further define $c_j^i = \max_t\{c_j^i(t)\}$ and we assume all the $t_{j,k}^i$ are i.i.d random variables for $1 \leq k \leq c_j^i$. In this model, the duplication of a particular task is made only once. Thus, the following equation holds:

$$t_j^i = \min\{(1-s_i)t_{j,1}^i, t_{j,2}^i, \ldots, t_{j,c_j^i}^i\} + s_i t_{j,1}^i \quad \forall i \tag{20}$$

In Section IV, we derive the distribution function of job *flowtime* by letting all the tasks of jobs being launched simultaneously. However, it is impossible to schedule the jobs in the same manner for the detection based model as stragglers can only be detected after the tasks have run for some time. In this way, it is difficult to simplify the expression for the job *flowtime*. As such, we choose to only optimize the total resource consumption of the job which yields the following formulation (P3):

$$\begin{aligned}
\min_{c_j^i(t), \sigma_i} \quad & \mathbb{E}\left[\sum_{i=1}\sum_{j=1}^{m_i} \gamma \cdot c_j^i \cdot t_j^i - \sum_{i=1}\sum_{j=1}^{m_i} \gamma \cdot (c_j^i - 1) \cdot s_i t_{j,1}^i\right] \\
s.t. \quad & \sum_{i=1}\sum_{j=1}^{m_i} c_j^i(t) \leq M \quad \forall t \\
& c_j^i = \max_t\{c_j^i(t)\} \quad \forall i; j \\
& 1 \leq c_j^i \leq r \quad \forall i; j \\
& (18), (19), (20)
\end{aligned}$$

Notice that P3 is a stochastic programming problem and we need to find the optimal solutions for $c_j^i(t)$ and σ_i.

A. Solving P3 through decomposition

The objective in this formulation can be decoupled as the summation of the expected resource consumption for each individual task. However, the first constraint makes this problem difficult to solve as it cannot be decoupled. We first relax this constraint and minimize the expected resource consumption for an individual task. The constraint will be taken into account later on when we design the actual scheduling algorithm.

Define $y_j^i = min\{t_{j,2}^i, t_{j,3}^i, \ldots, t_{j,c_j^i}^i\}$ if $c_j^i \geq 2$ and $y_j^i = \infty$ if $c_j^i = 1$. Let $d_j^i \triangleq min\{(1-s_i)t_{j,1}^i, y_j^i\}$. It can be readily shown that

$$\mathbb{E}\left[\gamma \cdot c_j^i t_j^i - \gamma \cdot (c_j^i - 1) s_i t_{j,1}^i\right] = \gamma \cdot \mathbb{E}\left[c_j^i d_j^i + s_i t_{j,1}^i\right] \tag{21}$$

Denote by π_j^i the event that a straggler is detected for task δ_j^i and ϑ_j^i the event that $y_j^i \leq \sigma_i \mathbb{E}[x_j^i]$ respectively. Then, we have

$$\begin{aligned}\mathbb{E}\left[c_j^i d_j^i + s_i t_{j,1}^i\right] = \mathbb{E}\left[c_j^i d_j^i + s_i t_{j,1}^i \left|\pi_j^i\right.\right] \cdot Pr\left(\pi_j^i\right) \\ + \mathbb{E}\left[c_j^i d_j^i + s_i t_{j,1}^i \left|\overline{\pi_j^i}\right.\right] \cdot Pr\left(\overline{\pi_j^i}\right)\end{aligned} \tag{22}$$

Further, denote by ϱ_j^i the event that $d_j^i > \sigma_i \mathbb{E}[x_j^i]$. Thus,

$$\mathbb{E}\left[c_j^i d_j^i \left|\pi_j^i\right.\right] = \mathbb{E}\left[c_j^i y_j^i \left|\vartheta_j^i\right.\right] \cdot Pr(\vartheta_j^i) + \mathbb{E}\left[c_j^i d_j^i \left|\varrho_j^i\right.\right] \cdot Pr(\overline{\vartheta_j^i}) \tag{23}$$

$$\mathbb{E}\left[c_j^i y_j^i \left|\vartheta_j^i\right.\right] \cdot Pr(\vartheta_j^i) = c_j^i \int_0^{\sigma_i \mathbb{E}[x_j^i]} t d(1-(1-F_i(t))^{c_j^i - 1}) \tag{24}$$

Define $G_i(t) = (1-F_i(t))^{c_j^i - 1}(1 - F_i(\frac{t}{1-s_i}))$, then

$$\mathbb{E}\left[c_j^i d_j^i \left|\varrho_j^i\right.\right] \cdot Pr(\overline{\vartheta_j^i}) = \frac{c_j^i}{1 - F_i(\frac{\sigma_i \mathbb{E}[x_j^i]}{1-s_i})} \int_{\sigma_i \mathbb{E}[x_j^i]}^{\infty} t d(1 - G_i(t)) \tag{25}$$

Combine Equality (23), (24), (25)

$$\begin{aligned}\mathbb{E}\left[c_j^i d_j^i \left|\pi_j^i\right.\right] = c_j^i \int_0^{\sigma_i \mathbb{E}[x_j^i]} (1-F_i(t))^{c_j^i - 1} dt \\ + \frac{c_j^i}{1 - F_i(\frac{\sigma_i \mathbb{E}[x_j^i]}{1-s_i})} \int_{\sigma_i \mathbb{E}[x_j^i]}^{\infty} (1 - G_i(t)) dt\end{aligned} \tag{26}$$

For a fixed σ_i, the optimal value of c_j^i is a function of σ_i and is determined by the following equation:

$$(c_j^i)^* = \rho_i(\sigma_i) = \arg\min_{c_j^i} \mathbb{E}\left[c_j^i d_j^i \left|\pi_j^i\right.\right] \tag{27}$$

And the optimal value of σ_i is determined by the following equation:

$$\sigma_i^* = \arg\min_{\sigma_i} \mathbb{E}\left[c_j^i d_j^i + s_i t_{j,1}^i \left|c_j^i = \rho_i(\sigma_i)\right.\right] \tag{28}$$

For the detailed steps, please refer to [17]. Define $\tau_i(c_j^i, \sigma_i) = \mathbb{E}\left[c_j^i d_j^i \left|\pi_j^i\right.\right]$. We conclude that, if $F_i(t)$ follows the pareto distribution, then $\tau_i(c_j^i, \sigma_i)$ is an increasing function of c_j^i when $c_j^i \in [2, \infty)$ for all $\alpha_i, \sigma_i > 1$ where α_i is the heavy-tail order. Moreover, we can choose an appropriate σ_i such that $\tau_i(2, \sigma_i) < \tau_i(1, \sigma_i)$. For the detailed proof, please refer to the technical report. Based on this conclusion, we derive the optimal solution of $c_i(t)$ for P3 in the following theorem:

Theorem 3. *The optimal value for $c_j^i(t)$ is 2 once a straggler is detected under the Pareto heavy-tail distribution. Moreover, the optimal value for σ_i does not depend on $\mathbb{E}[x_j^i]$ or s_i but the heavy-tail order of the Pareto distribution.*

Proof. Refer to [17]. □

B. The Design of the Straggler Detection Algorithm (SDA) based on P3

Based on the solution of P3, we propose the *Straggler Detection Algorithm* (SDA) to optimize the resource consumption for stragglers. SDA monitors the progress for each task running in the cluster. When a new job J_i arrives, the optimal value for c_j^i and σ_i are computed based on Equation (27), (28).

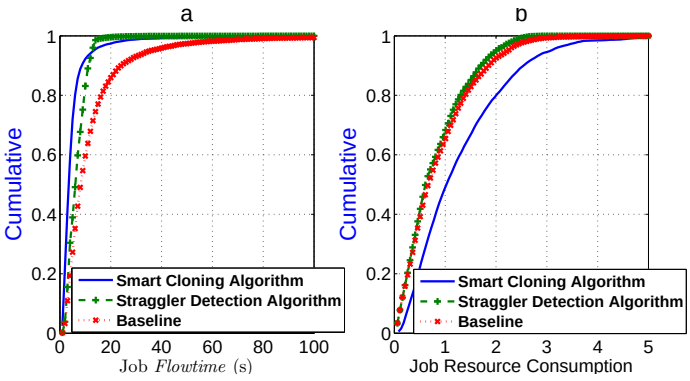

Fig. 2. The comparison between our proposed SCA and SDA, the Microsoft Mantri's speculative execution algorithm is adopted as the baseline. Panel a shows the cmf of Job *flowtime* and Panel b shows the cmf of total job resource consumption.

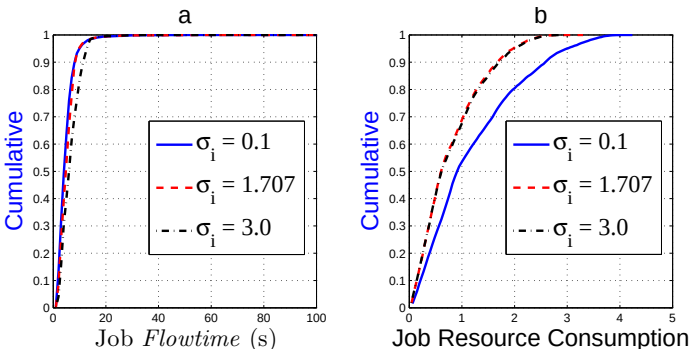

Fig. 3. The comparison between different σ_i in the SDA algorithm. It indicates that when $\alpha = 1.707$, both the job *flowtime* and resource consumption achieves the optimal.

SDA consists of three scheduling levels. In the first level, if the task δ_j^i is suffering from a straggler, the duplicates of δ_j^i are assigned to alternative machines. The policy is that for the straggler δ_j^i, the scheduler duplicates $c_j^i - 1$ new copies of it on other machines if resource is available and the machine is randomly chosen from any available ones. If the number of idle machine is less than $c_j^i - 1$, then all of these idle machines will be assigned a duplicate of δ_j^i. In the second level, the remaining tasks of the jobs which have begun but not finished yet are scheduled at the beginning of each time slot. The job with the smallest remaining workload has the highest priority. In the third level, the jobs have not begun their work yet will be scheduled in the cluster and among those which has the smallest workload will be scheduled first. Similarly, denote by $\chi(l) = \{J_{l_1}, J_{l_2}, \cdots\}$ the set of unscheduled jobs at time slot l and it's sorted in the same way as SCA. The scheduler assigns only one copy for each task in the newly scheduled job from $\chi(l)$.

C. Performance evaluation for SDA

We run another simulation to evaluate the performance of SDA. The simulation parameters are completely the same as SCA. For $\alpha = 2$, we apply Theorem 3 and Equation (28) to obtain the optimal value for σ_i which is $1 + \sqrt{2}/2$. The comparison between SDA and the baseline algorithm is shown

in Fig. 2. It shows that both the job *flowtime* and resource consumption of SDA perform better than the baseline though the objective of SDA is to optimize the resource consumption only. Fig. 3 shows the comparison result for different σ_i. It indicates that when $\sigma_i = 1+\sqrt{2}/2$, both the job *flowtime* and resource consumption achieve the best performance. Moreover, when σ_i decreases below $1+\sqrt{2}/2$, the total resource consumption increases. The job *flowtime* increases when σ_i increases above $1+\sqrt{2}/2$ because the speculative copy is made late and cannot help much for reducing the job *flowtime*. Compared to SDA, SCA performs better in terms of *flowtime* for small jobs. However, the average job *flowtime* of SCA is roughly the same while SCA consumes more resource. It is unfair to conclude that the SDA outperforms SCA. The reason is that SDA monitors the progress of each task and only makes the duplicates for the stragglers. However, SCA doesn't depend on the task progress and makes the cloning in a greedy fashion. Due to this reason, SCA tends to consume more resource than SDA. To determine which algorithm is better in real systems, the monitoring cost should be taken into account.

VI. Design of Straggler-Detection-based Algorithm for the Heavily Loaded Regime

In the heavily loaded cluster, i.e., $\lambda \geq \lambda^U$, the cloning-based scheme cannot be applied and only the Straggler-Detection-based approach is efficient. However, the SDA algorithm implemented in lightly loaded cluster launches the duplicates immediately when a straggler is detected. In the heavily loaded cluster, the resource is intensive and the scheduler should wait for the available machines. As a result, on-demand scheduling for stragglers is not always possible. In the literature, Microsoft Mantri [3] chooses to schedule a speculative copy if $\mathbb{P}(t_{rem} > 2 * t_{new}) > \delta$ is satisfied when there are available machines. We extend this scheme and propose the *Enhanced Speculative Execution* (ESE) algorithm. The scheduling decision is made in each time slot and we also leave some space for cloning the small jobs. Usually the small job represents for interactive applications like query and have very strict latency requirement. When the workload is heavy, the probability of making speculative copies for the tasks of small jobs is low in Mantri's scheduling algorithm. Cloning for small jobs only incurs a small amount of resource consumption while contributing a lot to the job *flowtime*.

A. The Enhanced Speculative Execution algorithm

The ESE algorithm also includes three scheduling levels. Explicitly, at the beginning of time slot l, the scheduler estimates the remaining time of each running task and puts the tasks whose remaining time satisfy a particular condition into the backup candidate set $D(l)$. Define $t^i_{j,rem}(l)$ as the remaining time of task δ^i_j at the beginning of time slot l. Then,

$$D(l) = \{\delta^i_j : c^i_j(l) = 1 \ \& \ t^i_{j,rem}(l) > \sigma_i \cdot \mathbb{E}[x^i_j]\}$$

All the tasks in $D(l)$ are sorted according to the decreasing order of $t^i_{j,rem}(l)$ and the scheduler assigns a duplicate of each task in $D(l)$ based on this order. In the same way as P3, we also need to find an appropriate value for σ_i that the expected resource consumption for a single task is minimized.

The scheduler then assigns the remaining tasks of the jobs which have already been scheduled but have not quitted the cluster yet. Denote by $R(l)$ the unfinished job set at time slot l and the jobs are sorted based on the remaining workloads. Upon scheduling, the jobs which have smaller remaining workload are given the higher priorities.

$N(l)$ is updated after the above scheduling. If there are still available machines and some jobs waiting to be scheduled, the scheduler then tries to do the cloning for small jobs and assign the number of copies which can maximize the difference between the job utility and total resource consumption. For big jobs, no cloning will be made. More precisely, denote by $\chi(l) = \{J_{l_1}, J_{l_2}, \cdots\}$ the job set which contains all the jobs have not been scheduled yet in the cluster. The jobs in $\chi(l)$ are sorted according to the increasing order of workloads. For small job J_{l_i} which satisfies $m_{l_i} < \eta \cdot \frac{N(l)}{|\chi(l)|}$ and $\mathbb{E}[x^{l_i}_j] < \xi$ in $\chi(l)$, the optimal number of copies cloned for task in J_{l_i} is determined by the following equation.

$$c^*_{l_i} = \arg\max_{c_{l_i}} U\left(\mathbb{E}[t_{l_i}], m_{l_i}\right) - \gamma \cdot \sum_{j=1}^{m_{l_i}} c_{l_i} \cdot \mathbb{E}[t^{l_i}_j] \qquad (29)$$

The parameters η and ξ can be tuned based on the cluster workload. Here, c_{l_i}, t_{l_i}, m_{l_i} and $t^{l_i}_j$ are the same as Section IV and $w^j_{m_{l_i}}$ is equal to l for all j.

The corresponding pseudocode is given in Algorithm 2 in the below.

B. Approximate Analysis for σ_i in ESE

σ_i has a great impact on the system performance. In the following analysis, we aim to find an appropriate σ_i in the ESE algorithm through minimizing the expected resource consumption for one particular task.

Define the random variable R^i_j as the resource that task δ^i_j consumes in the cluster. For convenience, we take $\gamma = 1$ and denote by θ^i_j the event that $x^i_j > \sigma_i \mathbb{E}[x^i_j]$. Then the expectation of R^i_j can be expressed in the following equation.

$$\mathbb{E}[R^i_j] = \mathbb{E}\left[R^i_j \left| \theta^i_j \right.\right] Pr(\theta^i_j) + \mathbb{E}\left[R^i_j \left| \overline{\theta^i_j} \right.\right] Pr(\overline{\theta^i_j}) \qquad (30)$$

Moreover, we have

$$\mathbb{E}\left[R^i_j \left| \overline{\theta^i_j} \right.\right] Pr(\overline{\theta^i_j}) = \int_0^{\sigma_i \mathbb{E}[x^i_j]} t d(F_i(t)) \qquad (31)$$

We give the following definition which helps to derive the expression of $\mathbb{E}[R^i_j]$.

Definition 2. *The asktime of a running task i is the earliest time that the scheduler checks whether it should duplicate a new copy for task i on available machines or not.*

Assume the interval of each time slot is short enough, then the *asktime* of δ^i_j can be treated as uniformly distributed in the interval [0, x^i_j] due to the cluster is heavily loaded. At this

Input: The jobs in the cluster associated with their running status at time slot l;
Output: Scheduling decisions for time slot l.

Count $N(l)$, the number of idle machines at time slot l and update $D(l)$, $R(l)$, $\chi(l)$.
for *the task* δ_j^i *in in* $D(l)$ **do**
  Assign a duplicate of δ_j^i on a random idle machine;
  $N(l)$ -= 1;
  if $N(l) == 0$ **then**
    return;
  end
end
for *the job* J_i *in* $R(l)$ **do**
  Assign the unscheduled tasks of J_i on idle machines;
  update $N(l)$;
  if $N(l) == 0$ **then**
    return;
  end
end
for *the job* J_{l_i} *in* $\chi(l)$ **do**
  if $m_{l_i} < \eta \cdot \frac{N(l)}{|\chi(l)|}$ & $\mathbb{E}[x_j^{l_i}] < \xi$ **then**
    Compute $(c_{l_i})^*$ base on Equation (29);
    Assign $(c_{l_i})^*$ duplicates for each task in J_{l_i};
    update $N(l)$;
  else
    Assign one duplicate for each task in J_{l_i};
    update $N(l)$;
  end
  if $N(l) == 0$ **then**
    return;
  end
end
return;

Algorithm 2: Enhanced Speculative Execution Algorithm

asktime, the scheduler assigns a new copy for δ_j^i if $t_{rem} > \sigma_i \mathbb{E}[t_{new}]$ is satisfied. Hence,

$$\mathbb{E}\left[R_j^i \,|\theta_j^i\right] Pr(\theta_j^i) = \int_{\sigma_i \mathbb{E}[x_j^i]}^{\infty} d(F_i(t))[\int_0^{t-\sigma_i \mathbb{E}[x_j^i]} \frac{1}{t}(x + 2\mathbb{E}[\min\{t-x, t_{new}\}])dx + \sigma_i \mathbb{E}[x_j^i]] \tag{32}$$

In Equation (32), $\mathbb{E}[min\{t-x, t_{new}\}]$ is given by:

$$\mathbb{E}[min\{t-x, t_{new}\}] = \int_0^{t-x} w d(F_i(w)) + (t-x)(1-F_i(t-x)) \tag{33}$$

Combining Equation (30),(31),(32),(33), $\mathbb{E}[R_j^i]$ can be determined and it's a function of σ_i. We obtain the optimal value of σ_i that minimizes the expected resource consumption through letting the derivative of $\mathbb{E}[R_j^i]$ be 0, .

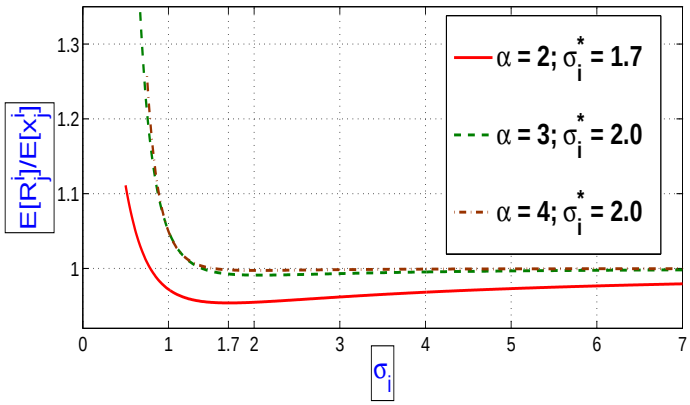

Fig. 4. The illustration of $E[R_j^i]/E[x_j^i]$ under different σ_i for pareto distribution when $\alpha = 2, 3, 4, 5$.

We illustrate the picture of $\mathbb{E}[R_j^i]$ in Fig. 4 under different σ_i for the pareto distribution. α is the heavy-tail order and σ_i^* is the optimal value. It indicates that when σ_i is close to 1.7 where α is 2, $\mathbb{E}[R_j^i]$ achieves the minimum value. We also compare the different optimal value for σ_i when the heavy-tail order changes. As shown in Fig. 4, σ_i^* increases along with α. Moreover, for all $\alpha \geq 3$, σ_i^* is very close to 2.0.

C. Performance Evaluation for ESE

In this subsection, we first conduct one simulation to show the impact of σ_i and the heavy-tail order on the cluster performance to demonstrate the fitness of our approximate analysis. Then we continue to compare the performance of ESE algorithm and Mantri's speculative execution strategy.

1) The impact of σ_i and α: In this simulation, there is only one job which consists of 10000 tasks and the cluster has $M = 100$ computing nodes. The expected task duration is 1 and the other parameters are as same as the simulation for SCA. For each $\sigma_i \in (0, 6)$ we run 50 simulations and take the average. We adopt the naive scheme in which speculative execution is not implemented as a comparison to the ESE algorithm. The simulation results are illustrated in Fig. 5. It shows that when σ_i is close to 1.7 for $\alpha = 2$, both the total amount of resource consumption and job *flowtime* achieve the minimum. Next, we keep the same expected task duration and take $\alpha = 2, 3, 4$ to show the impact of heavy-tail order on the optimal value of σ_i. The result matches well with the theoretical analysis in Fig. 4. It also indicates that as α increases, the performance improvement of the ESE algorithm tends to be inconspicuous for a single job case.

2) The performance of ESE Algorithm: In this simulation, there are $M = 3000$ machines in the cluster. The simulation parameters are the the same as the SCA and SDA algorithms except for the job arrival rate. Here, we adopt two arrival rates which are 30 and 40 jobs per unit time. We follow the approximate analysis and choose σ_i to be 1.7. η and ξ are set to be 0.1 and 1 respectively. Similarly, we use the Microsoft Mantri's algorithm as the baseline to compare it with the ESE algorithm. The results are illustrated in Fig. 6. As shown, when the job arrival rate is 40, 80% of jobs can

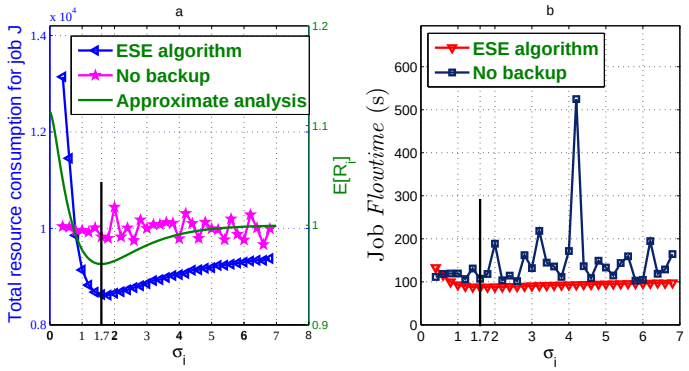

Fig. 5. Simulation result for a single job. Panel a shows the comparison of resource consumption under different σ_i between ESE algorithm and the naive method without backup. The blue line represents the ESE algorithm, the dark green line describes the result of the approximate analysis while the pink red line represents the method without backup. Panel b shows the comparison of job *flowtime* time.

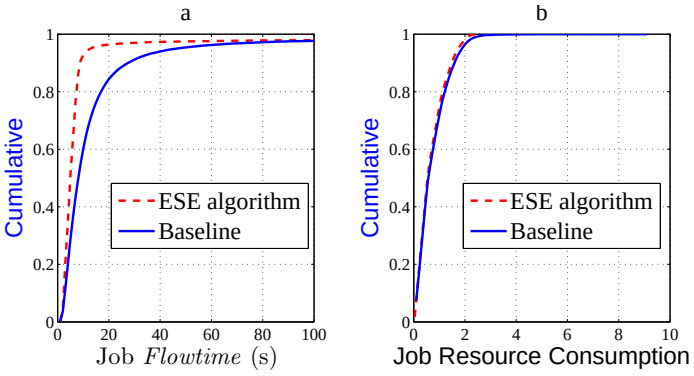

Fig. 6. The comparison between our proposed ESE Algorithm and the Microsoft Mantri's speculative execution algorithm. The job arrival rate is 40. Panel a shows the cmf of Job *flowtime* and Panel b shows the cmf of total job resource consumption.

finish within 10 units time in the ESE algorithm while 80% of jobs can finish only within 18 units time in the baseline. Further, we observe that the average job *flowtime* reduces by 18% in the ESE algorithm compared to the baseline. However, the resource consumption for these two algorithms are roughly the same. This is because cloning for small jobs can incur an extra amount of resource consumption. When the job arrival rate λ is 30, the job resource consumption for ESE algorithm reduces a lot compared to the baseline. We also apply the SCA and SDA algorithms to the heavily loaded cluster but the performance turns out to be poor. As the cluster is heavily loaded, the cloning in SCA easily blocks the scheduling of newly arriving jobs. Due to resource is intensive, there is a little chance to make on-demand scheduling for stragglers in SDA and thus can not help to improve the performance.

VII. Conclusions

In this paper, we address the speculative execution issue in the parallel computing cluster and focus on two metrics which are job *flowtime* and resource consumption. We utilize the distribution information of task duration and build a theoretical framework for making speculative copies. We categorize the cluster into lightly loaded and heavily loaded cases and derive the cutoff threshold for these two operating regimes. Moreover, we propose two different strategies when the cluster is lightly loaded and design the ESE algorithm while the cluster is heavily loaded. This is the first work so far to adopt the optimization-based approach for speculative execution in a multiple-job cluster. In the future work, we will consider to make speculative execution for the cluster where there exists task dependency in a job. Like the MapReduce framework, any reduce task can only begin after the map tasks finish within a job.

References

[1] Apache. http://hadoop.apache.org, 2013.

[2] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica. Effective straggler mitigation: Attack of the clones. In *the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2013.

[3] G. Ananthanarayanan, S. Kandula, A. Greenberg, I. Stoic, Y. Lu, B. Saha, and E. Harris. Reining in the outliers in mapreduce clusters using mantri. In *USENIX OSDI*, Vancouver, Canada, October 2010.

[4] D. Breitgand, R. Cohen, A. Nahir, and D. Raz. On cost-aware monitoring for self-adaptive load sharing. *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS*, 28(1):70–83, January 2010.

[5] H. Chang, M. Kodialam, R. R. Kompella, T. V. Lakshman, M. Lee, and S. Mukherjee. Scheduling in mapreduce-like systems for fast completion time. In *Proceedings of IEEE Infocom*, pages 3074–3082, March 2011.

[6] F. Chen, M. Kodialam, and T. Lakshman. Joint scheduling of processing and shuffle phases in mapreduce systems. In *Proceedings of IEEE Infocom*, March 2012.

[7] Q. Chen, C. Liu, and Z. Xiao. Improving mapreduce performance using smart speculative execution strategy. *IEEE Transactions on Computers*, PP(99), January 2013.

[8] R. B. Cooper. *Introduction to queueing theory*. The Macmillan Company, New York, 1972.

[9] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *Proceedings of OSDI*, pages 137–150, December 2004.

[10] M. Isard, M. Budiu, Y. Yu, A. Birrell, and D. Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. In *Proceeding of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems*, March 2007.

[11] Lefechetz, S., Lasalle, and J.P. *Stability by Lyapunov's direct method with applications*. Academic Press, New York, 1961.

[12] B. Moseley, A. Dasgupta, R. Kumar, and T. Sarlos. On scheduling in map-reduce and flow-shops. In *Proceedings of SPAA*, pages 289–298, June 2011.

[13] X. Sun, C. He, and Y. Lu. Esamr: An enhanced self-adaptive mapreduce scheduling algorithm. In *the 18th International Conference on Parallel and Distributed Systems (ICPADS)*, December 2012.

[14] J. Tan, X. Meng, and L. Zhang. Delay tails in mapreduce scheduling. In *Proceedings of SIGMETRICS*, pages 5–16, London, United Kingdom, June 2012.

[15] W. Wang, K. Zhu, L. Ying, J. Tan, and L. Zhang. Map task scheduling in mapreduce with data locality: Throughput and heavy-traffic optimality. In *Proceedings of IEEE Infocom*, Turin, Italy, April 2013.

[16] Q. Xie and Y. Lu. Degree-guided map-reduce task assignment with data locality constraint. In *Information Theory Proceedings (ISIT)*, pages 1–6, July 2012.

[17] H. Xu and W. C. Lau. Optimization for speculative execution of multiple jobs in a mapreduce-like cluster. In *Tech. Rep*, http://personal.ie.cuhk.edu.hk/%7Exh112/Research/technical-report.pdf, February 2014.

[18] M. Zaharia, A. Konwinski, A. D. Joseph, R. Katz, and I. Stoica. Improving mapreduce performance in heterogeneous environments. In *Proceeding of the 8th USENIX conference on Operating systems design and implementation*, December 2008.

[19] Y. Zheng, N. Shroff, and P. Sinha. A new analytical technique for designing provably efficient mapreduce schedulers. In *Proceedings of IEEE Infocom*, Turin, Italy, April 2013.